

CLAIMS

What is claimed is:

1. A method of transmitting data frames to a plurality of output ports, each of the data frames having a destination associated with one of the output ports, the method comprising:

5 at each of a plurality of input ports, partitioning a portion of each data frame to provide one or more ordered data cells having data representative of a sequence number corresponding with the output port associated with ^athe destination of the data frame, the data representative of the sequence number in each data cell indicating an ordinal position of the data cell among the ordered data cells of the data frame; and

10 at each of the output ports, receiving a forwarded data cell for each ordered data cell associated with each data frame having a destination associated with the output port, each forwarded data cell corresponding with an ordered data cell and data frame associated with the ordered data cell, and determining an ordinal position of the forwarded data cell among the forwarded data cells associated with the data frame based upon data in the forwarded data cell
15 representative of the sequence number.

2. The method of claim 1, the method further comprising, at each input port:

maintaining a sequence number corresponding to each output port;

transmitting the ordered data cells of data frames having a destination associated with the

20 output port sequentially, one ordered data cell at a time; and

incrementing the sequence number for each subsequent ordered data cell transmitted to the output port from the input port.

3. The method of claim 2, the method further comprising:

at a crossbar, receiving ordered data cells transmitted from each of the input ports and enqueueing up to a maximum number of received ordered data cells of data frames having a destination associated with any one particular output port; and

5 at each input port, restarting the sequence number for each output port following a transmission of at least the maximum number of ordered data cells of data frames that can arrive out of order having a destination associated with the output port.

4. The method of claim 1, the method further comprising, for at least one input port:

10 receiving a first data frame having a destination associated with a particular one of the output ports, the first data frame having a priority for transmission to the particular output port; transmitting ordered data cells of the first data frame to a crossbar;

while transmitting the ordered data cells of the first data frame, receiving a second data frame having a destination associated with the particular output port and having a priority for transmission to the particular output port higher than the priority for transmitting the first data frame to the particular output port; and

15

completing the transmission of each ordered data cell of the first data frame to the crossbar prior to a transmission of any ordered data cell of the second data frame to the crossbar.

20 5. The method of claim 1, wherein each forwarded data cell received at the output ports includes source information identifying the input port providing the ordered data cell associated with the forwarded data cell, the method further including at the output port:

associating each received forwarded data cell with an input port based upon the source information in the received forwarded data cell;

for each received forwarded data cell associated with an input port, determining an ordinal position of the received forwarded data cell among forwarded data cells associated with a data frame based upon the sequence number in the received forwarded data cell.

6. The method of claim 1, wherein each forwarded data cell received at the output ports includes source information identifying the input port providing the ordered data cell associated with the forwarded data cell, the method further including at the output port:

associating each received forwarded data cell with an input port based upon the source information in the forwarded data cell;

for each received forwarded data cell associated with an input port, determining an ordinal position of the received forwarded data cell among forwarded data cells associated with a data frame based upon the sequence number in the received forwarded data cell.

7. A data switch comprising:

a plurality of output ports for transmitting forwarded data frames to destinations;
a plurality of input ports for receiving data frames, each received data frame having a destination associated with one of the output ports, each of the plurality of input ports including logic for partitioning a portion of each received data frame to provide one or more ordered data cells having data representative of a sequence number corresponding with the output port associated with the destination of the received data frame, the data representative of the sequence number in

each ordered data cell indicating an ordinal position of the ordered data cell among the ordered data cells of the data frame,

wherein each of the output ports receives forwarded data cells, each forwarded data cell corresponding with an ordered data cell generated at one of the input ports and having data indicative of the sequence number of the corresponding ordered data cell, and includes logic for determining an ordinal position of the forwarded data cell among the forwarded data cells of a forwarded data frame based upon the data indicative of the sequence number in the forwarded data cell.

10 8. The data switch of claim 7, wherein each input further includes:

logic for maintaining a sequence number corresponding to each output port;

logic for scheduling transmission of the ordered data cells of received data frames having a destination associated with the output port sequentially, one data cell at a time; and

15 logic for incrementing the sequence number for each subsequent ordered data cell transmitted to the output port from the input port.

9. The data switch of claim 8, wherein each of the ordered data cells includes a data payload having a portion of a data payload of a corresponding received data frame, and wherein the data switch further comprises a crossbar for receiving ordered data cells from the input ports and, for
20 each received ordered data cell, transmitting a forwarded data cell having the data payload of the received ordered data cell to an output port associated with the destination of the received data frame corresponding with the received ordered data cell.

10. The data switch of claim 9, wherein the crossbar is capable of buffering up to a maximum number of received ordered data cells of received data frames having a destination associated with any one particular output port, and wherein each input port further includes logic for restarting the sequence number for each output port following the transfer of at the maximum
5 number of ordered data cells of received data frames which can arrive out of order having a destination associated with the output port.

11. The data switch of claim 7, wherein each input port further includes:

logic for scheduling transmission of ordered data cells of a first received data frame to a
10 particular output port associated with the destination of the first received data frame, the first received data frame having a priority for transmission to the particular output port; and

logic for completing the transmission of each of the ordered data cells of the first data frame to the output port prior to a transmission of any ordered data cell of a second received data frame to the output port, the second received data frame having a destination associated with the
15 output port and the second received data frame having a priority for transmission to the particular output port higher than the priority for transmitting the first data frame to the particular output port.

12. The data switch of claim 7, wherein each outbound data cell received at the output ports
20 includes source information identifying the input port transmitting the associated data cell to the crossbar, and the output ports further include:

logic for associating each received outbound data cell with an input port based upon the source information in the received outbound data cell; and

for each received outbound data cell associated with an input port, logic for determining an ordinal position of the received outbound data cell among outbound data cells associated with a data frame based upon the sequence number in the received data outbound data cell.

- 5 13. In a data communication network including a plurality of host computers for transmitting data packets to a plurality of network devices, each of the data packets having data representative of a destination network address, each of the network devices having a media access control (MAC) address associated therewith, the improvement including:

10 a plurality of output ports, each of the output ports being coupled to at least an associated one of the network devices for transmitting MAC data frames to the at least one network device according to the MAC address associated therewith;

a look-up engine for receiving the data packets from the host computers addressed to one or more of the network devices and forming intermediate data frames based upon the data packets, the intermediate data frames having a data payload and information identifying an
15 output port associated with the one or more network devices;

a plurality of input ports for receiving intermediate data frames from the look up engine, each received data frame having a destination associated with one of the output ports, each of the plurality of input ports including logic for partitioning the data payload of each received intermediate data frame to provide one or more ordered data cells having data representative of a
20 sequence number corresponding with the output port associated with the destination of the received intermediate data frame, the data representative of the sequence number in each ordered data cell indicating an ordinal position of the ordered data cell among the ordered data cells of the intermediate data frame,

wherein each of the output ports receives forwarded data cells, each forwarded data cell corresponding with an ordered data cell originating at one of the input ports and having data indicative of the sequence number of the corresponding ordered data cell, and includes logic for determining an ordinal position of the forwarded data cell among the forwarded data cells of a forwarded data frame based upon the data indicative of the sequence number in the forwarded data cell.

14. The data communication network of claim 13, wherein each input port further includes:
logic for maintaining a sequence number corresponding to each output port;

logic for scheduling transmission of the ordered data cells of received intermediate data frames having a destination associated with the output port sequentially, one data cell at a time; and

logic for incrementing the sequence number for each subsequent ordered data cell transmitted to the output port from the input port.

15. The data communication network of claim 14, wherein each of the ordered data cells includes a data payload having a portion of the data payload of a corresponding received intermediate data frame, and wherein the improvement further comprises a crossbar for receiving ordered data cells from the input ports and, for each received ordered data cell, transmitting a forwarded data cell having the data payload of the received ordered data cell to an output port associated with the destination of the received intermediate data frame corresponding with the received ordered data cell.

16. The data communication network of claim 15, wherein the crossbar is capable of enqueueing up to a maximum number of received ordered data cells of received intermediate data frames having a destination associated with any one particular output port, and wherein each input port further includes logic for restarting the sequence number for each output port following the transfer of at least the maximum number of ordered data cells of received intermediate data frames that can arrive out of order having a destination associated with the output port.

17. The data communication network of claim 13, wherein each input port further includes:

logic for scheduling transmission of ordered data cells of a first received intermediate data frame to a particular output port associated with the destination of the first received intermediate data frame, the first received intermediate data frame having a priority for transmission to the particular output port; and

logic for completing the transmission of each of the ordered data cells of the first received intermediate data frame to the output port prior to a transmission of any ordered data cell of a second received intermediate data frame to the output port, the second received intermediate data frame having a destination associated with the output port and the second received intermediate data frame having a priority for transmission to the particular output port higher than the priority for transmitting the first received intermediate data frame to the

particular output port.

18. The data communication network of claim 13, wherein each forwarded data cell received at an output port includes source information identifying the input port providing the ordered data cell associated with the forwarded data cell, and wherein the output ports further include:

logic for associating each received forwarded data cell with an input port based upon the
5 source information in the received forwarded data cell; and

for each received forwarded data cell associated with an input port, logic for determining an ordinal position of the received forwarded data cell among forwarded data cells associated with a data frame based upon the sequence number in the received forwarded data cell.